



An Oracle White Paper
August 2011

Oracle VM 3: Architecture and Technical Overview

Introduction	1
Overview of Oracle VM 3 Concepts and Components.....	3
Key Deployment Concepts.....	4
Oracle VM Manager	4
Centralized Network and Storage Configuration and Management	5
Dynamic Resource Management.....	6
Oracle VM Server for x86.....	7
Hypervisor and Management Domain.....	7
Cluster File System – OCFS2	8
Oracle VM Agent.....	8
Types of Virtual Machines	8
Oracle VM Templates and Virtual Assemblies for Applications	9
Conclusion	11
Appendix 1: References.....	12

Introduction

Oracle VM is an enterprise-class server virtualization solution comprised of [Oracle VM Server for x86](#) and Oracle VM Manager. Oracle VM 3 reflects Oracle's strategic commitment to deliver Application Driven Virtualization—Virtualization that makes the entire enterprise software and hardware stack easier to deploy, manage, and support so IT and business can be more agile.

Oracle VM has been designed from the ground-up for excellent scalability, manageability, and ease-of-use. Oracle VM 3 delivers significant enhancements to make it faster and easier to rollout operating systems, enterprise application and middleware across your *entire* data center—not just the Oracle software environment—while reducing costs and making your datacenter or cloud environment highly available and secure.

[Oracle VM Server for x86](#) is the most scalable x86 server virtualization solution in the market today, and it has been tested to handle mission critical enterprise workloads with support for up to 160 physical CPUs and 2 TB of memory. For virtual machines, Oracle VM 3 can support up to 128 virtual CPUs and 1TB memory per guest VM. Oracle VM supports industry standard x86 operating systems and servers from Oracle and other leading vendors, and it supports a broad range of network and storage devices, making it easy to integrate into your environment.

[Oracle VM Manager](#) provides an easy-use-centralized management environment for configuring and operating your server, network, and storage infrastructure from a browser-based interface (no Java client required), and it is accessible from just about anywhere. Users leverage the virtualization manager for creating management policies, as well as cloning, sharing, configuring, booting and migrating VMs. Oracle VM helps customers improve server utilization, achieve higher availability and achieve better performance while reducing costs.

Beyond the virtualization layer, Oracle offers solutions such as [Oracle Virtual Assembly Builder](#), [Oracle VM Templates](#), [Oracle Optimized Solution for Enterprise Cloud Infrastructure](#) that leverage Oracle VM and provide application-driven virtualization. With these tools you can:

- Quickly configure and provision entire multi-tier enterprise applications onto virtualized and cloud environments
- Deploy a fully configured software stack with preinstalled and preconfigured software images
- Save application deployment time by up to 98 percent

Oracle VM 3 supports easy integration with [Oracle Enterprise Manager](#) to provide comprehensive management of the full hardware and software stack from applications to disk. Management of the virtual environment can be handled with the many automated features of Oracle VM Manager or a combination of Oracle VM Manager and Oracle Enterprise Manager (see figure 1).

Oracle VM Server for x86 is the only x86 server virtualization solution fully certified by Oracle, allowing for the best support and highest availability for the complete software stack. Oracle products such as Oracle Database, Oracle Fusion Middleware, and Oracle Applications are officially certified when running in Oracle VM-based environments, and they have all gone through vigorous real-world testing. Customers who obtain subscription service for Oracle VM receive access to patches, fixes, and updates via [Unbreakable Linux Network \(ULN\)](#) and 24x7 global support.

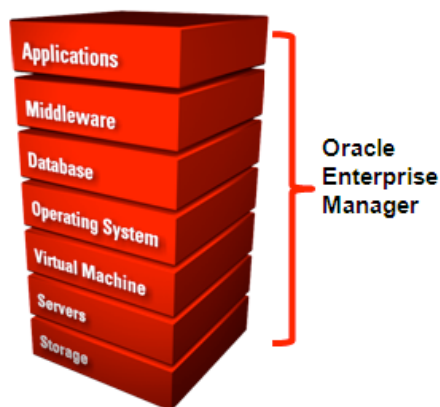


Figure 1. Oracle Full Stack Integration

Overview of Oracle VM 3 Concepts and Components

Running on x86 servers, Oracle VM includes Oracle VM Manager and Oracle VM Server for x86.

- Oracle VM Manager controls the virtualization environment, creating and monitoring Oracle VM servers and the virtual machines. Oracle VM Manager serves as the only administrative interface to the Oracle VM servers, unlike previous-generations of Oracle VM that were jointly administered from the management server, as well as locally from the command-line for each Oracle VM Server. Oracle VM Manager 3 is an Oracle Fusion Middleware application, based on the Oracle Weblogic Server application server and Oracle Database. The Oracle VM Manager runs on 64-bit Oracle Linux 5.5 OS or later. For its management repository, Oracle VM Manager uses an Oracle Database, which can be installed either on the same management server or a separate server. Oracle Database 11g Express Edition (XE) is bundled with Oracle VM Manager 3 for customer evaluation purposes, but Oracle Database Standard Edition (SE) or Enterprise Edition (EE) is required for production support as the management repository. All necessary licenses, including licenses for Weblogic Server and Database (SE or EE), are included at no additional charge.
- Oracle VM Server for x86 installs directly on server hardware with x86 Intel or AMD processors and does not require a host operating system. An Oracle VM Server is comprised of a hypervisor and privileged domain (Dom0) that allows multiple domains or virtual machines (i.e. Linux, Solaris, Windows, etc.) to run on one physical machine. The Dom0 runs a process called Oracle VM Agent. The Oracle VM Agent receives and processes management requests, provides event notifications and configuration data to the Oracle VM Manager. Oracle VM Server 3 requires 64-bit x86 hardware, but can support either 64-bit or 32-bit guest virtual machines.

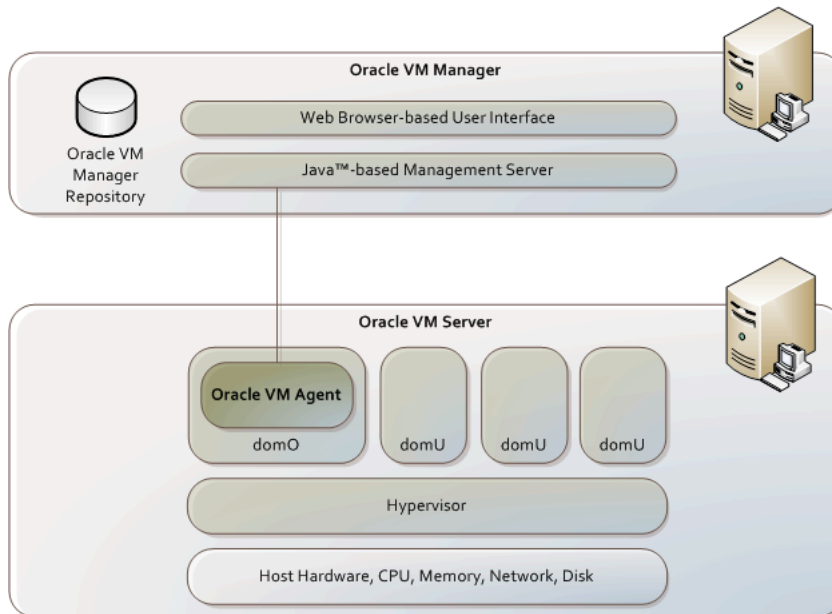


Figure 2. Oracle VM Components

Key Deployment Concepts

Multiple Oracle VM Servers are grouped into server pools as shown in Figure 3. Each server pool can have up to 32 physical servers, and every server in a given pool has access to shared storage, which can be NFS, Fibre Channel or iSCSI (or a combination of these). This allows VMs associated with the pool to start and run on any physical server within the pool. *Note:* Local storage can also be configured but is often not appropriate for production environment since it prevents or sharply constrains the ability for a VM to run anywhere in the pool in the event of a server failure or live migration event.

When a VM first starts in the server pool, the VM is placed on the server that is the “best” fit based on an algorithm that takes into account the server CPU, memory and network utilization. Typically the server that is assigned to host the VM has the most resources available. Then the VM is associated with the server, but it can be securely migrated to any other server in the server pool by the administrator or can be load balanced across the pool based on a dynamic resource scheduling policy.

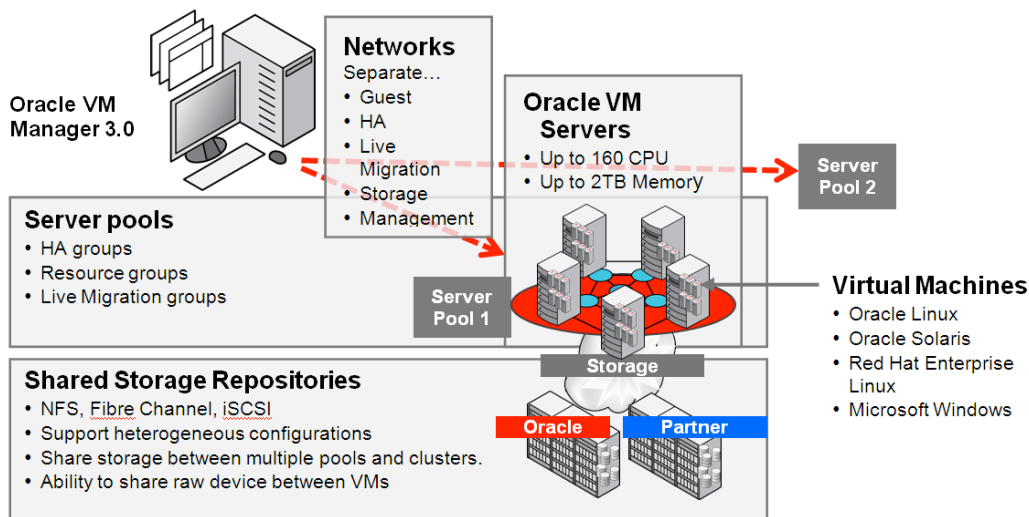


Figure 3. Oracle VM Deployment

As a result of this architecture, VMs can easily start-up, power-off, migrate, and/or restart without being blocked by the failure of any individual server or by the failure of multiple servers as long as there are adequate resources in the pool to support the requirements for all VMs to run concurrently.

Oracle VM Manager

Oracle VM Manager has been significantly enhanced in version 3 to simplify administrative operations, provide completely centralized network and storage management, and improve dynamic resource management for the virtualization environment.

Oracle VM Manager 3 is built with Oracle Weblogic Server, providing enterprise-grade transaction management for operations, including historical event tracking, and the ability to cleanly abort long-running operations. Every Oracle VM Manager operation is performed as a structured task known as a

“job” in Oracle VM Manager. A job consists of all the steps involved in the successful completion of the operation. The information provided for each job can be viewed from the GUI, and includes status, progress (percent complete and which steps have completed), and the steps remaining. It is possible to abort a job at any time, or if the job is aborted by the system because of a terminal failure, the state is rolled back to its original state. Jobs provide a simple way to see who made changes within the environment and trace failures.

Oracle VM Manager 3 uses Java and XML within the Oracle ADF 11g (Application Development Framework). ADF generates JavaScripts to accommodate various browser-rendering. Asynchronous JavaScript and XML (AJAX) and it is used to give the UI a richer look by providing a way to update aspects of the UI asynchronously and without requiring the user to hit the browser refresh button to get an updated view. The Oracle VM Manager also includes a fully interactive tree view and automatic refresh to show the real-time state of both the physical and virtual environment.

Oracle VM Manager 3 provides the only true browser-based user interfaces compared to other enterprise server virtualization vendors that rely on Java clients to be installed. Oracle VM Manager’s browser interface makes it easier to access Oracle VM Manager from nearly anywhere on nearly any device supporting standard browsers.

Centralized Network and Storage Configuration and Management

Network and storage are key parts of any application deployment. Oracle VM simplifies these steps by centralizing all configurations through the Oracle VM Manager and by creating pool-level network settings and logical storage to reduce administrative time.

Network Configuration and Management

In order to provide the highest and most predictable performance, scalability, and availability for the VMs, Oracle VM is designed to support sophisticated network configurations. The networking infrastructure in the Oracle VM environment needs to be able to support a number of functions. For example:

- The guest network supporting IP traffic between Oracle VM Servers and the outside world
- The management network between Oracle VM Servers and Oracle VM Manager
- The storage network between the Oracle VM Servers and their storage sub-systems

Oracle VM provides the ability to separate networks in order to isolate traffic and maintain a predictable quality of service. Up to five separate network types can be defined from Oracle VM Manager: A server management network, live migration network, cluster heartbeat network, virtual machine IP network, and a storage network.

When you create an Oracle VM network, you map available network ports to a set of logical Ethernet networks from the Oracle VM Manager GUI. All Oracle VM Server logical network configuration and management tasks—such as creating networks, setting up NIC port bonding, and configuring VLAN Networks—are now performed using Oracle VM Manager.

Storage Configuration and Management

The Oracle Storage Connect framework enables Oracle VM Manager to directly leverage the resources and functionality of existing storage systems in the Oracle VM environment, supporting native storage services such as SAN or NFS storage creation, deletion, and expansion. This allows the Oracle VM Manager to automatically discover available storage using Oracle VM Manager and create new storage repositories or map raw storage directly to virtual machines. To learn more, please refer to the [Oracle Storage Connect white paper](#).

Dynamic Resource Management

Automating the virtualized environment to improve application quality of service and reduce power consumption are key enablers for improving operational efficiency.




Figure 4. Policy Based Dynamic Resource Scheduling

Enhancements in Oracle VM 3 include the following:

- **Distributed Resource Scheduling (DRS) for capacity management.** DRS provides real-time monitoring of Oracle VM Server utilization with the goal to rebalance a server pool to provide consistent resources to the running virtual machines. DRS migrates load away from heavily loaded Oracle VM Servers to servers running lighter loads. (See Figure 4)
- **Dynamic Power Management (DPM) to optimize server pool for minimum power consumption.** DPM complements DRS to reduce the servers in the pool when there are periods of low resource utilization. It can automatically add capacity as needed when resource utilization ramps up. (See Figure 5)

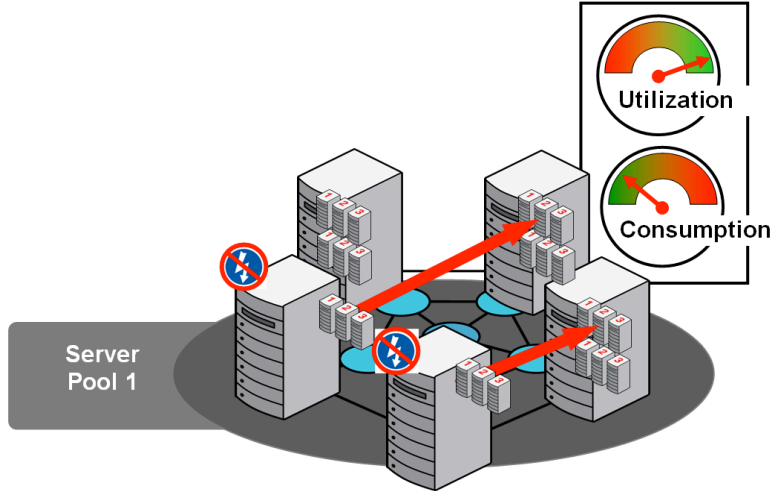


Figure 5. Policy Based Dynamic Power Management

Oracle VM Server for x86

Oracle VM Server 3 is installed on bare metal 64-bit x86 servers and supports 32-bit and 64-bit guest operating systems such as Linux, Solaris, and Windows. Oracle VM Server incorporates an open source Xen hypervisor component, which has been customized and optimized to integrate into the larger, Oracle-developed virtualization server. Oracle VM Server includes Oracle VM Agent to communicate with Oracle VM Manager for management and configuration. While Oracle VM Server software uses several open source components and is itself open source, it is a product developed and offered only by Oracle and that is completely unique to Oracle.

Hypervisor and Management Domain

The hypervisor runs directly on the hardware serving as the abstraction layer for all hardware operations from the guest OS such as CPU, I/O, and disk requests. By separating the guest VMs from the hardware, the hypervisor is able to run multiple operating systems securely and independently.

There's a privileged management domain running on the hypervisor with direct hardware access and guest management responsibilities. This management domain is known as Domain0 (Dom0). Hardware support (such as support for the network and storage devices) is provided by the device drivers installed in Dom0.

The Dom0 has unique privileges to access the hypervisor that are not allocated to any other guest domains, which are known as User Domains (DomUs). These privileges allow it to manage all aspects of guest domains such as starting, stopping, I/O requests, etc. Users can perform all the management functions from the easy-to-use Oracle VM Manager 3 interface, and there's no need to log into the Dom0 to perform guest management or device configuration.

The Dom0 of Oracle VM Server 3 is specially built to handle enterprise-class workload requirements. It leverages Oracle Linux with the 64-bit Oracle Unbreakable Enterprise Kernel (UEK), but has been

customized and minimized for tight security and scalability in a virtualized environment. The Oracle Unbreakable Enterprise Kernel is a fast, modern, reliable Linux kernel that is optimized for enterprise software and hardware.

The optimization of the hypervisor and special-built Oracle Linux with Unbreakable Enterprise Kernel has enabled Oracle VM Server the most scalable x86 server virtualization solution. For example, Oracle VM Server 3.0 supports up to 160 physical CPUs and 2TB memory, and each guest VM supports up to 128 virtual CPUs (vCPUs) and 1TB memory. Such scalability limits have been tested on [Oracle's Sun Fire X4800 M2 Server](#).

Cluster File System – OCFS2

[OCFS2](#) (Oracle Cluster File System 2) is a [free, open source](#), general purpose, extent-based clustered file system that Oracle developed and contributed to the Linux community. OCFS2 provides an open source, enterprise-class alternative to proprietary cluster file systems, and provides both high performance and high availability.

OCFS2 1.8 is incorporated into Oracle VM Server 3 as the underlying cluster file system for server pools as well as the storage repository. In addition, advanced features including thin provisioning and instant clone in OCFS2 enable significantly faster and more efficient virtual machine provisioning and cloning.

Oracle VM Agent

The Oracle VM Agent runs on the Dom0 and handles requests from the Oracle VM Manager. It processes requests by interacting with the Oracle VM hypervisor, Oracle VM Storage Connect plug-in, the Distributed Lock Manager (DLM), and the HA framework provided by OCFS2.

When integrated with Oracle Enterprise Manager, Oracle's system management offering for managing the complete application stack, top-to-bottom, the Oracle VM 3 agent also serves as a proxy agent for Enterprise Manager to eliminate the need to deploy additional management agents.

Types of Virtual Machines

Guest virtual machines running on Oracle VM Server can be configured to use one of the following virtualization modes:

- **Paravirtualized (PVM).** In paravirtualized mode, the kernel of the guest operating system is modified to distinguish that it is running on a hypervisor instead of on the bare metal hardware. As a result, I/O actions and system clock timers in particular are handled more efficiently, as compared with non-paravirtualized systems where I/O hardware and timers have to be emulated in the operating system. Oracle VM supports PV kernels for Oracle Linux and Red Hat Enterprise Linux, offering better performance and scalability.
- **Hardware virtualized or fully virtualized (HVM).** When support for hardware virtualization (Intel VT-x or AMD-V) is available in the host server hardware, the guest operating system may run

unmodified and depend on the hardware features to ensure proper operation in a virtualized environment. However, HVM can introduce additional overhead since it relies on binary translation and device emulation. But it brings the benefits of compatibility: largely any “unmodified” operating system image such as Linux, Solaris, or Windows that is not virtualization-aware can run as an HVM guest without any changes, making it easy to consolidate existing physical server images, leaving them essentially untouched. With the recent development in Intel and AMD processors, in some scenarios, hardware-assisted virtualization will also provide performance benefits.

- **Hardware virtualized using paravirtualized drivers (PVHVM).** This mode is identical to hardware virtualized (HVM), but with additional paravirtualized I/O drivers installed in the guest's operating system to improve virtual machine performance. This allows the operating system to essentially run unmodified, but uses PV drivers that have been modified specifically to execute I/O commands efficiently in a virtualized environment to minimize overhead. Oracle provides PV drivers for Microsoft Windows. PV drivers are provided as part of the operating system for Oracle Linux and Oracle Solaris.

For details of comparing the different virtualization modes, please refer to My Oracle Support Knowledge Note ([Doc ID 757719.1](#)): Comparison of Guest Virtualization Modes, HVM, PVM and PVHVM.

Oracle Linux with the Unbreakable Enterprise Kernel can be run directly on bare metal or as a virtual guest on Oracle VM.

Tip: The Unbreakable Enterprise Kernel version kernel-uek-2.6.32-100.34.1 added support for paravirtualized drivers in a HVM guest on Oracle VM. Starting with this kernel version, the default is to present only paravirtualized drivers when running in a hardware-virtualized guest. To run kernel-uek – including the drivers – fully hardware-virtualized, an additional kernel boot parameter “xen_emul_unplug=never” must be added to the boot parameters in `/etc/grub.conf`.

Oracle VM Templates and Virtual Assemblies for Applications

Oracle VM provides the ability to rapidly and easily deploy a single pre-built, pre-configured, pre-optimized, and pre-patched guest virtual machine (or multiple machines depending on the application). The guest VM can contain a complete Oracle software solution along with the operating system and related software infrastructure. These guest VMs, called Oracle VM Templates, are available from Oracle’s E-Delivery website and are ready to download and deploy. Already configured for production use, Oracle VM Templates can save users days or weeks learning to install and configure a sophisticated product such as Siebel CRM or Oracle Enterprise Manager Grid Control. Instead, users can simply download and start the VM to begin using it right away.

Within these Templates, Oracle software is automatically configured in the same manner as if the software had been installed and patched manually. The exact same directories and Oracle “homes” are used, and the package and patch inventories are completely standard and up-to-date so that no changes to operational procedures are required to maintain the instances over time. Accordingly, Oracle VM Templates can also be fully customized post-install and then re-saved as “golden image” Templates in

Oracle VM. Such Templates can serve as a user’s enterprise deployment standard to minimize risks and variation across multiple instance deployments.

All of the existing Oracle VM Templates that were certified with Oracle VM server 2.x are also certified with Oracle VM Server 3. An up-to-date list of available Templates can be found on Oracle VM Templates website at www.oracle.com/technetwork/server-storage/vm/templates-101937.html.

An application “assembly” is very similar to an Oracle VM Template except that an additional set of configuration information and management policies are packaged along with the set of multiple virtual machines, their virtual disks, and the inter-connectivity between them. These assemblies are packaged using the industry-standard Open Virtualization Format (OVF) for the configuration and policy information. Assemblies can be created as a set of .ovf Open Virtualization Format and disk image (.img) files, or may all be contained in a single Open Virtualization Format Archive (.ova) file.

Oracle Virtual Assembly Builder is a sophisticated development tool for introspecting the current application environment and creating assemblies or collections of VMs along with all their critical configuration parameters to allow you to move them to a virtual environment or to replicate easily and quickly within an existing environment.

The Oracle VM Guest VM API allows the guest to set/get key/value pairs through the Oracle VM Manager to the virtual machine. The API allows direct integration between guest software and the virtualization layer, to assist in orchestration of complex, multi-VM deployments.

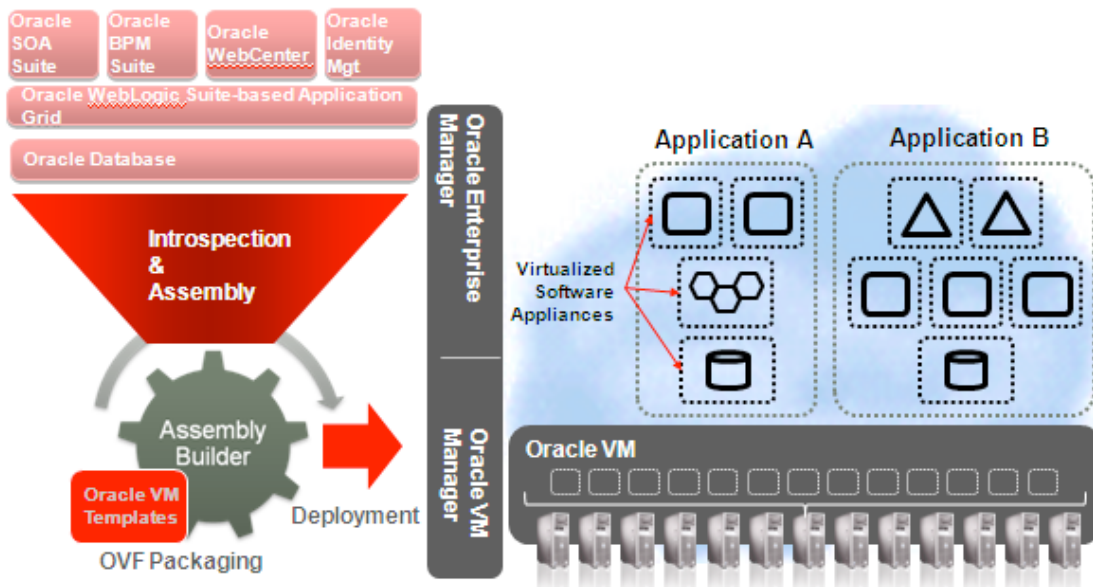


Figure 6. Oracle Virtual Assembly Builder

Conclusion

Successful businesses are defined by efficiency and speed. Those companies that use capital more efficiently and just-in-time generate higher profits and have more to invest in future initiatives. IT is a critical component of these organizations' competitive advantage. The back-office functions that provide critical services can no longer afford to be a rigid and balkanized organization without optimal efficiency. Application-driven virtualization is a way to improve IT's speed, efficiency and agility.

By deploying cloud computing, IT is responding to business requirements for increased agility while at the same time reducing costs and risks from the IT balance sheet. IT becomes much more responsive to the business, and through this alignment, a part of the business process.

By virtualizing with Oracle VM Server for x86 and taking advantage of unique Oracle tools, businesses can reduce their total cost of ownership in the datacenter, increase IT flexibility, and achieve greater business agility.

Oracle VM 3 demonstrates Oracle's commitment to deliver application driven server virtualization solutions. The new capabilities introduced in Oracle VM provide increased levels of scalability, manageability, and ease-of-use to help customers make the most demanding enterprise applications easier to deploy, manage, and support.

The best way to experience all the benefits of Oracle VM 3 is to download the software from Oracle's E-Delivery site at: <http://edelivery.oracle.com/oraclevm>, and try it out in your environment.

For more information about Oracle's virtualization, visit www.oracle.com/virtualization.

Appendix 1: References

For more information, visit the Web resources listed in Table 1.

TABLE 1. WEB RESOURCES FOR FURTHER INFORMATION	
WEB RESOURCE DESCRIPTION	WEB RESOURCE URL
Oracle Virtualization	http://oracle.com/virtualization
Oracle VM Templates	http://www.oracle.com/technetwork/server-storage/vm/templates-101937.html
Oracle Solaris	http://oracle.com/solaris
Oracle Linux	http://oracle.com/linux
Sun Servers	http://www.oracle.com/servers
Sun Storage	http://www.oracle.com/storage
Sun Networking	http://www.oracle.com/us/products/servers-storage/networking/
Oracle Enterprise Manager	http://www.oracle.com/us/products/enterprise-manager/
Oracle Enterprise Manager Ops Center	http://www.oracle.com/us/products/enterprise-manager/opscenter/



Oracle VM 3:
Architecture and Technical Overview
August 2011

Author: Honglin Su
Contributing Authors: Adam Hawley, Susan
Roberts, Chris Kawalek

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together